

From Exergy to Intelligence:

A General Theory of Economic Transduction

Matthew Long

The YonedaAI Collaboration

YonedaAI Research Collective

Chicago, IL

matthew@yonedaai.com · <https://yonedaai.com>

23 July 2026

Abstract

Economic production often begins with energy and ends with an outcome that is described in a different language: useful work, a communicated message, a verified task, an adopted decision, or a change in welfare. Treating those objects as if they were one conserved substance invites category errors. This paper gives a typed accounting framework for studying such chains without reducing their later stages to thermodynamic state variables. A transduction graph has nodes that carry a semantic quantity type, a physical or accounting unit, and a declared system boundary. Directed stages carry observed input and output amounts. Along a path whose interfaces match, the terminal to initial ratio equals the product of the stage yields. The result is a telescoping identity, not a law that converts joules into intelligence or money. We prove the identity, its covariance under changes of reporting unit, a pathwise stochastic version, and failure results for mismatched boundaries, duplicated outputs, and products of expectations.

The physical part of the framework uses standard definitions of exergy and useful work. The informational part uses Shannon entropy and mutual information. Landauer's bound is stated only for logically irreversible operations under a thermal model. The economic part introduces task verification, adoption, counterfactual attribution, and a value functional as additional typed stages. Information, intelligence, institutions, and value are not thermodynamic state variables. Three bounded numerical examples show what the accounting identity can establish and what it leaves open. A standard-library Python implementation validates typed interfaces and reproduces deterministic and Monte Carlo decompositions. The strongest result is therefore narrow but useful: heterogeneous production stages can share an auditable graph when their interfaces are explicit, while the graph itself refuses unsupported conversions.

Keywords: exergy; useful work; information theory; thermodynamics of computation; production networks; causal attribution; economic measurement; compositional accounting
Evidence statement. The path composition theorem and unit covariance proposition are PROVED under their stated assumptions. Numerical examples and software tests are

COMPUTATIONAL. Claims about realized economic value are **CONDITIONAL** on a named counterfactual and identification design. A universal physical measure of intelligence or value remains **OBSTRUCTED**.

1 Introduction

An electric motor, a communication channel, an inference server, and a court registry all use energy. That common fact does not make their outputs the same kind of object. Shaft work has units of joules. A decoded message can be described in bits relative to a probability model. A correct answer is defined by a task and scoring rule. A court judgment changes rights only within an institutional order. An economic benefit depends on a counterfactual, an affected population, and a value functional. Moving from one description to the next requires an explicit interface.

The phrase “energy becomes intelligence” can therefore mean two very different things. In a modest engineering sense, an energized device executes operations that participate in a system producing answers. In a stronger and usually mistaken sense, energy is imagined as a substance that persists as information, intelligence, and value. The first claim can be measured. The second has no basis in thermodynamics, information theory, or economics. Energy is conserved in a closed physical accounting. Exergy is destroyed by irreversibility. Shannon information is defined by probability distributions. Economic value is not conserved at all.

This paper develops a formal language for the modest claim. The main object is a typed transduction graph

$$E \longrightarrow X \longrightarrow O \longrightarrow U \longrightarrow A \longrightarrow \Delta W,$$

where the letters stand for declared quantities rather than universal substances. Depending on the application, E may be facility electricity, X may be exergy or a named computation count, O may be an immediate output, U a verified useful outcome, A an adopted action, and ΔW a counterfactual change in a stated value functional. Each arrow has its own numerator, denominator, boundary, and evidence source.

For a linear path with matched interfaces, the familiar product

$$\frac{\Delta W}{E} = \frac{X}{E} \frac{O}{X} \frac{U}{O} \frac{A}{U} \frac{\Delta W}{A}$$

is an algebraic identity. The intermediate quantities cancel because the denominator of one factor is the same measured quantity as the numerator of the preceding factor. This condition is stricter than sharing a label. A facility energy measurement cannot silently become device energy. A generated answer count cannot silently become a verified answer count. An observed revenue change cannot silently become a causal welfare effect.

The identity is useful precisely because it is not profound. It localizes losses, missing evidence, and category changes. If a model produces many tokens but few correct task outcomes, the output to outcome stage is weak. If correct recommendations are ignored, the adoption stage is weak. If an adopted recommendation has no identified counterfactual effect, the final stage is unknown. The graph lets a researcher say where the evidence stops.

The contribution has four parts. First, we define typed nodes and composable stages that include semantic type, unit, boundary, time interval, and attribution metadata. Second, we prove deterministic and stochastic composition results together with failure conditions. Third, we place exergy, useful work, Shannon information, and Landauer’s principle in their proper domains. Fourth, we give bounded examples and an executable reference model.

The novelty claim is deliberately limited. Exergy analysis, useful-work economics, input-output models, production networks, thermoeconomics, information theory, causal graphs, and compositional systems theory all predate this work. The new proposal is a common audit discipline for chains that cross those literatures while preserving their type boundaries. Whether that discipline proves empirically useful is an open question.

1.1 Questions and answers

The paper addresses five questions.

1. What must be recorded before two adjacent production stages can be multiplied?
2. Which statements about exergy and useful work remain physical?
3. Which statements about bits and erasures remain informational or thermodynamic?
4. Where do task usefulness, adoption, institutions, and economic value enter?
5. What can a product of stage yields establish in a stochastic or branching system?

The short answers are: interfaces must be identical, exergy is relative to an environment, bits are relative to a probability model, economic stages require task and causal definitions, and naive multiplication fails when interfaces, boundaries, allocations, or random variables do not match.

2 Related work and the boundary of the contribution

2.1 Exergy, useful work, and thermoeconomics

Classical thermodynamics distinguishes energy quantity from work potential. Exergy analysis formalizes the maximum useful work obtainable as a system comes to equilibrium with a specified environment [32, 19, 6, 12]. Exergy destruction measures irreversibility through the Gouy-Stodola relation. It depends on an ambient temperature and, for chemical exergy, a reference composition. These dependencies are not nuisances. They are part of the definition.

Thermoeconomics and exergoeconomics attach monetary costs to exergy streams, components, or products to diagnose production systems [33, 34]. Such methods join a physical balance to a cost allocation rule. The currency figures do not become thermodynamic state variables merely because they appear beside exergy in one model.

Economy-wide useful-work research estimates the conversion from primary or final exergy to delivered end-use services and studies its relation to growth. Ayres and Warr [3], Warr et al. [35], and Warr and Ayres [36] are direct predecessors of the present program. Cleveland et al. [10], Stern [30], and Kümmel [20] likewise treat energy as an important production input. This paper does not claim to discover that relation.

The useful-work literature also gives a warning. National energy series, sectoral final energy, useful mechanical work, and energy services are different aggregates. Their construction requires conversion efficiencies, allocation conventions, and historical data. A general transduction graph should expose those bridge assumptions rather than placing every series under one symbol E .

2.2 Information and thermodynamics

Shannon [29] defined entropy and channel capacity for communication under a probability model. His theory abstracts from meaning. Landauer [21] linked logically irreversible information processing to a minimum heat cost in a thermal environment. Bennett [7, 8] showed why logical reversibility matters and why the cost cannot be assigned indiscriminately to every computational step. Modern stochastic thermodynamics makes the relation between information and physical work more precise while preserving model assumptions [25].

Research on predictive and semantic information asks when correlations are useful to an agent or support viability [31, 18]. These models add a task, agent, or viability functional. They do not recover meaning from the Shannon measure alone. That distinction is central here.

2.3 Economic production and networks

Leontief input-output analysis records interindustry flows [22]. Hulten's growth accounting result [14] and modern production-network work [1, 9] show how shocks and productivity changes propagate through linked sectors. Those models typically work with commodities, prices, or production functions rather than semantic type systems, but they establish that graph structure matters economically.

The present graph is closer to an accounting and measurement graph than a general equilibrium model. It does not solve prices or assert market clearing. It can be embedded in those models, but it first asks a prior question: are the quantities on adjacent edges actually the same measured interface?

2.4 Compositional systems and resource theories

Open systems, Petri nets, string diagrams, and resource theories provide formal accounts of composition [5, 4, 11, 13]. A typed interface prevents a component from being connected to an incompatible component. Our construction borrows that discipline but uses elementary directed acyclic graphs so that economic and engineering readers can inspect the accounting without categorical machinery.

The distinction matters for novelty. The theorem proved below is a telescoping identity. Its value lies in the metadata required before the telescope is allowed. The claim is not that graph composition itself is new.

2.5 Causal attribution and welfare

Potential-outcome and graphical causal models distinguish an observed outcome from a counterfactual effect [27, 26, 15]. Social choice and welfare economics show that aggregation requires preferences or weights [2, 28]. These results constrain the final stage of an economic transduction chain. Measured output can be physical. Realized value is defined only relative to a counterfactual and evaluative rule.

Remark 2.1 (Novelty boundary). The framework unifies reporting requirements across existing fields. It does not replace thermodynamics, information theory, production-network economics, or causal inference. A successful application must still use the substantive model appropriate to each stage.

3 Measurement primitives

3.1 The measurement tuple

Every reported terminal-to-energy ratio is relative to a measurement tuple

$$\mathcal{M} = (b, c, W, \tau, e, a, \pi),$$

where b is the physical and organizational boundary, c is the counterfactual, W is a value or outcome functional, τ is the time horizon, e is the energy convention, a is the attribution rule, and π is the uncertainty model. This tuple follows the program-wide notation. A transduction graph refines it by recording local metadata at every interface.

Definition 3.1 (Quantity type). A quantity type is a tuple

$$\theta = (q, d, b, \tau, c, a),$$

where q is a semantic quantity name, d is a unit dimension, and the remaining elements declare boundary, interval, counterfactual when relevant, and attribution rule when relevant. Two types are equal only if every component needed for the comparison is equal.

The semantic name is not redundant. Electrical energy and exergy may both be reported in joules, yet they are not the same quantity. A generated response and a verified response may both be counts, yet the latter has passed a scoring rule. Gross revenue and incremental welfare may both be reported in currency, yet their counterfactual meanings differ.

Definition 3.2 (Typed amount). A typed amount is a pair (x, θ) with $x \in \mathbb{R}_{\geq 0}$ and quantity type θ . The number x is a magnitude in the reporting unit declared by θ .

The non-negativity convention suits flow magnitudes. Net economic effects can be negative. We represent them with a signed terminal amount or, when stage yields must remain non-negative, with separate benefit and harm channels. Neither representation should be chosen silently.

3.2 Boundaries

A boundary is more than a box around hardware. At minimum it names:

1. the physical equipment and infrastructure;
2. the start and end times;
3. the geographic and grid location;
4. the energy stage, such as device, wall, facility, final, primary, or lifecycle energy;
5. treatment of idle operation, retries, failures, and shared capacity;
6. allocation of embodied energy and common infrastructure;
7. the functional unit and evaluated population;
8. exclusions.

These declarations block common substitutions. Thermal design power is not measured electricity. Device telemetry is not facility energy. Facility electricity is not lifecycle energy. A primary-energy series cannot be added to the final energy produced from that primary input. Such errors are not small changes in efficiency. They change the denominator.

3.3 Semantic types and dimensions

Let $\mathcal{D}$ be the free abelian group generated by reporting units. Examples of generators are J, bit, task, and USD₂₀₂₆. Quotient units live in $\mathcal{D}$, so a stage may have unit task/J. Dimensional cancellation is necessary for composition, but it is not sufficient. Semantic type and boundary must also match.

This separation answers a recurring objection. A count of bits and a count of tasks are dimensionless in some conventions, but that does not make them interchangeable. The graph carries both an algebraic unit and a semantic type. The type checker rejects a bit-to-task substitution unless a named evaluation stage connects them.

Definition 3.3 (Evidence record). An evidence record for a typed amount contains the observation procedure, instrument or data source, calibration or validation method, uncertainty, version, and responsible boundary declaration. A modeled value must be marked as modeled rather than observed.

3.4 Stocks, flows, and services

The framework distinguishes a stock from a flow and a service. Energy is often reported as a flow over an interval. Installed compute capacity is a stock. Completed tasks are event counts. A reduction in waiting time is a service outcome. A welfare change can extend beyond the measurement interval. Dividing any of these by joules is syntactically possible, but comparability requires matching temporal support and discounting.

Power and energy must also remain separate:

$$P(t) = \frac{dE(t)}{dt}, \quad E[t_0, t_1] = \int_{t_0}^{t_1} P(t) dt.$$

A ratio per watt is not automatically a ratio per joule. The former may describe a rate under steady operation; the latter describes an integrated quantity over a declared interval.

4 Typed transduction graphs

Definition 4.1 (Transduction graph). A typed transduction graph is a finite directed graph

$$\mathcal{G} = (N, L, s, t, \Theta, x),$$

where N is a set of nodes, L a set of directed stages, $s, t : L \rightarrow N$ are source and target maps, $\Theta : N \rightarrow \text{Type}$ assigns a quantity type to each node, and x assigns non-negative measured or modeled amounts to stage interfaces.

Each stage $\ell : i \rightarrow j$ has an input amount x_ℓ^{in} of type $\Theta(i)$ and an output amount x_ℓ^{out} of type $\Theta(j)$. For positive input, its numerical yield is

$$y_\ell = \frac{x_\ell^{\text{out}}}{x_\ell^{\text{in}}},$$

with quotient unit d_j/d_i .

Definition 4.2 (Composable path). A directed path $\mathcal{P} = (\ell_1, \dots, \ell_n)$ is accounting-composable when:

1. $t(\ell_k) = s(\ell_{k+1})$ for every adjacent pair;
2. the shared semantic type and metadata are identical;
3. $x_{\ell_k}^{\text{out}} = x_{\ell_{k+1}}^{\text{in}}$ in a common reporting unit;
4. no shared output is counted twice through an undeclared branch.

The third requirement can use a declared tolerance for measured data. A nonzero residual then appears as a reconciliation item, not as hidden efficiency loss.

4.1 A canonical chain

Figure 1 gives the canonical sequence used throughout the paper. The labels below the arrows are quotient types, not universal conversion constants.

$$E_{\text{wall}} \xrightarrow{y_1} C_{\text{named}} \xrightarrow{y_2} O_{\text{immediate}} \xrightarrow{y_3} U_{\text{verified}} \xrightarrow{y_4} A_{\text{adopted}} \xrightarrow{y_5} \Delta W_c$$

Figure 1: A typed path from wall energy to counterfactual value. Every node requires its own operational definition and boundary.

An application may stop at any node. A hardware benchmark can responsibly report C/E_{wall} . A task benchmark can report U/E_{wall} if it observes the intervening stages. It need not claim $\Delta W/E$. Refusing to fill an unsupported edge is a feature of the model.

4.2 Acyclicity and feedback

The reference implementation requires a directed acyclic graph for a single accounting interval. Real economies contain feedback. Prices affect adoption, adoption affects capacity, and learning changes later efficiency. We represent feedback by unrolling time:

$$A_t \rightarrow D_{t+1} \rightarrow E_{t+1} \rightarrow O_{t+1}.$$

This preserves acyclicity at the event level. A simultaneous equilibrium model can sit outside the graph and supply stage values, but it must not be mistaken for a directly observed chain.

4.3 Losses and residuals

For physical conversion, a loss may be an energy or exergy flow. For task evaluation, a failed output is not thermodynamic waste in the same sense. For adoption, nonuse may be rational under local information. The general word “loss” should therefore be read as a stage shortfall relative to a declared output, not as entropy production.

Let r_k denote a reconciliation residual at interface k :

$$x_{\ell_k}^{\text{out}} = x_{\ell_{k+1}}^{\text{in}} + r_k.$$

If $r_k \neq 0$, the simple product theorem below does not apply without a bridge stage that accounts for r_k .

5 Composition results

Theorem 5.1 (Stage-yield composition). *Let $\mathcal{P} = (\ell_1, \dots, \ell_n)$ be an accounting-composable path with positive initial input. Then*

$$\frac{x_{\ell_n}^{\text{out}}}{x_{\ell_1}^{\text{in}}} = \prod_{k=1}^n y_{\ell_k}.$$

The product unit is the terminal unit divided by the initial unit.

Proof. By definition,

$$\prod_{k=1}^n y_{\ell_k} = \frac{x_{\ell_1}^{\text{out}}}{x_{\ell_1}^{\text{in}}} \frac{x_{\ell_2}^{\text{out}}}{x_{\ell_2}^{\text{in}}} \dots \frac{x_{\ell_n}^{\text{out}}}{x_{\ell_n}^{\text{in}}}.$$

Composability gives $x_{\ell_k}^{\text{out}} = x_{\ell_{k+1}}^{\text{in}}$ for $k = 1, \dots, n-1$. Every intermediate numerical amount cancels, leaving

$$\prod_{k=1}^n y_{\ell_k} = \frac{x_{\ell_n}^{\text{out}}}{x_{\ell_1}^{\text{in}}}.$$

The unit factors cancel by the same adjacent equalities. Semantic equality at each interface licenses those cancellations. Without it, identical algebraic symbols could refer to different quantities. $\square$

Remark 5.2 (Status and interpretation). Theorem 5.1 is PROVED. It is an accounting identity under exact interface matching. It does not show that the first quantity causes the last, that any stage is efficient in a welfare sense, or that the yield remains stable outside the observed interval.

Proposition 5.3 (Unit covariance). *Suppose the numerical representation of node i changes by a positive scale factor α_i . Then each yield transforms as*

$$y'_k = \frac{\alpha_{k+1}}{\alpha_k} y_k,$$

and the path product transforms as

$$\prod_{k=1}^n y'_k = \frac{\alpha_n}{\alpha_0} \prod_{k=1}^n y_k.$$

Thus the product agrees with the directly transformed terminal-to-initial ratio.

Proof. Substitute $x'_i = \alpha_i x_i$ into every adjacent ratio. Intermediate scale factors cancel. Only the terminal and initial factors remain. $\square$

The proposition explains why a stage yield can change numerically when joules become kilowatt-hours or dollars become cents without changing the measured relation. A published scalar without its unit is not invariant.

Proposition 5.4 (Pathwise stochastic composition). *Let $X_0, \dots, X_n$ be positive random variables on one probability space, with pathwise shared interfaces. Define $Y_k = X_k/X_{k-1}$. Then*

$$\frac{X_n}{X_0} = \prod_{k=1}^n Y_k$$

almost surely.

Proof. For every sample point on which all X_k are positive, the intermediate random variables cancel algebraically. The equality therefore holds almost surely. $\square$

Limitation 5.5 (Products of expectations). In general,

$$\mathbb{E}\left[\frac{X_n}{X_0}\right] \neq \prod_{k=1}^n \mathbb{E}[Y_k].$$

Equality needs additional conditions. Independence of the Y_k is sufficient for the expectation of their product to factor when the relevant moments exist, but stage yields in production systems are often dependent.

Dependence is substantively plausible. High load can raise hardware throughput while reducing task quality through batching or latency. Harder tasks can use more energy and fail more often. Selection into adoption can depend on expected benefit. Monte Carlo propagation must state the joint uncertainty model.

Proposition 5.6 (Mismatch obstruction). *If two adjacent stages report different interface amounts or different quantity types, no exact terminal-to-initial factorization follows from their reported yields alone.*

Proof. Let the first stage end at amount u and the second begin at amount v . The product contains the factor u/v :

$$\frac{u}{x_0} \frac{x_2}{v} = \frac{x_2}{x_0} \frac{u}{v}.$$

It equals the direct ratio only when $u = v$. If the types differ, the quotient u/v is not licensed even when their numerical representations happen to match. $\square$

5.1 Elasticities do not telescope in the same way

Average yields telescope because they are ratios of shared levels. Marginal responses require derivatives:

$$\frac{dX_n}{dX_0} = \prod_{k=1}^n \frac{dX_k}{dX_{k-1}}$$

only for a differentiable deterministic chain with no omitted common causes and with each X_k locally a function of X_{k-1} . In a network with multiple inputs, the total derivative includes all active paths. Average value per joule must not be read as marginal value of the next joule.

6 Branching, joint production, and attribution

Linear chains are the easiest case. Economic systems branch, merge, reuse intermediate outputs, and produce joint products. A graph can represent those features, but the scalar product theorem applies to one reconciled path at a time.

6.1 Splits

Suppose output X is allocated to two branches with shares a_1, a_2 . Physical conservation may justify

$$a_1 + a_2 \leq 1$$

for a material or energy flow. The unallocated remainder is then recorded. For nonrival information, copying can send one message to both branches without dividing its content. An allocation share is still needed if the researcher attributes shared energy or cost to each use, but that share is an accounting convention rather than a conservation law for information.

Definition 6.1 (Attribution-complete split). A split is attribution-complete when every unit of a rival input is assigned once, common infrastructure has a published allocation rule, and replicated nonrival outputs are marked as replicated rather than physically multiplied.

This definition prevents a common double count. If one model response informs two decisions, summing the gross benefits of both decisions may be valid only if their counterfactual effects are separate and nonoverlapping. The energy denominator cannot be charged in full to each branch and then recombined as if the branches used distinct energy.

6.2 Merges and complements

A stage with inputs $X_1, \dots, X_m$ is described by a production relation

$$Z = f(X_1, \dots, X_m; \omega),$$

where ω includes technology and context. No unique scalar yield from one input follows unless the other inputs and the attribution rule are fixed. If the inputs are complements, a zero value for one may collapse the output.

For such a stage, one can report:

1. the full vector of inputs per output;
2. a cost share or exergy allocation under a named rule;
3. a marginal product at a named operating point;
4. a Shapley or other cooperative allocation for a stated value function;
5. a sensitivity range across admissible allocations.

None of these is a free consequence of graph topology. The choice can change rankings.

6.3 Joint products and waste

Industrial processes often yield a principal product, coproducts, emissions, and waste heat. An AI service can yield a correct answer, a reusable cache entry, and a harmful error. An

institution can settle a dispute while creating precedent that affects later cases. A scalar terminal quantity suppresses this vector.

Hallucinated, toxic, or otherwise harmful model outputs belong on an outcome liability branch. They are not thermodynamic entropy. Their incidence and harm need a scoring rule, affected population, and counterfactual, just like beneficial task outcomes.

Let the terminal outcome be

$$\mathbf{Z} = (Z_{\text{benefit}}, Z_{\text{harm}}, Z_{\text{option}}, Z_{\text{externality}}).$$

A scalar $W(\mathbf{Z})$ requires weights, prices, preferences, or a welfare functional. The graph should publish $\mathbf{Z}$ before aggregation. A negative component cannot be hidden by reporting only gross benefit.

6.4 Cycles through time

Learning and institutional adaptation create cycles. A recommendation changes behavior; behavior changes future training data; the updated system changes later recommendations. Unrolling the graph by date creates a sequence

$$O_t \rightarrow A_t \rightarrow D_{t+1} \rightarrow O_{t+1}.$$

The energy and value denominators must then specify whether later gains are credited to the original intervention, the update, or both. Discounting and depreciation become explicit stages.

Limitation 6.2 (No conservation of value). The sum of economic values at outgoing branches need not equal the value of the incoming quantity. Value can be created, destroyed, transferred, or reassigned when preferences, scarcity, rights, or counterfactuals change. This is one reason value is not a thermodynamic state variable.

7 The physical layers: energy, exergy, and useful work

7.1 Energy and exergy

For a closed system relative to an environment at temperature T_0 and pressure p_0 , a common expression for nonflow exergy is

$$B = (U - U_0) + p_0(V - V_0) - T_0(S - S_0) + E_{\text{kin}} + E_{\text{pot}},$$

with chemical terms added when composition differs from the environment [19, 12]. The expression is defined relative to the reference environment. Changing T_0 , p_0 , or environmental composition can change the exergy.

For a control volume, exergy accounting separates fuel exergy, product exergy, destruction, and losses. In a common steady-state notation,

$$B_F = B_P + B_D + B_L, \quad B_D = T_0 S_{\text{gen}} \geq 0.$$

The exergy efficiency is

$$\eta_{\text{ex}} = \frac{B_P}{B_F},$$

where “fuel” and “product” are functional exergy definitions, not necessarily literal fuel and merchandise.

Energy and exergy answer different questions. The first law balances energy. The second law limits conversion and identifies irreversibility. Electricity has high work potential under common reference conditions. Low-temperature heat can carry substantial energy but little exergy. A model that uses only energy can miss quality differences among physical carriers.

7.2 Useful physical work

Useful work is the mechanical work or equivalent service delivered at an end use under a declared convention. Economy-wide studies reconstruct useful work by combining final energy by carrier and device with estimated conversion efficiencies [3, 35]. The resulting series is not directly observed in every sector. It inherits uncertainty from historical energy balances, device stocks, load factors, and conversion assumptions.

In a motor and pump chain, for example,

$$E_{\text{electric}} \rightarrow W_{\text{shaft}} \rightarrow W_{\text{hydraulic}} \rightarrow W_{\text{delivered}}.$$

Every node can be expressed in joules, and physical losses can be reconciled. The delivered service, such as liters of potable water at a specified head, adds another functional type. Economic benefit from that water requires prices, scarcity, health effects, distribution, and a counterfactual.

7.3 Why exergy does not extend automatically

Exergy is sometimes used metaphorically for social capacity, knowledge, or economic potential. The metaphor should not be used as a measurement identity. To be an exergy quantity,

a variable must belong to a thermodynamic model with a state, an environment, and a physically defined maximum work process. Institutional legitimacy and task correctness do not meet that definition.

Proposition 7.1 (No type-free exergy extension). *Equality of reporting units does not imply equality of quantity types. In particular, assigning joule units to a modeled social variable does not make that variable exergy.*

Proof. Under the type definition, equality requires semantic identity, boundary, and reference conditions in addition to the algebraic unit. A social score labeled in joule equivalents lacks the thermodynamic state and environment needed for the exergy type unless an explicit physical mapping is supplied. Unit equality alone is therefore insufficient. $\square$

7.4 Embodied and operational energy

Lifecycle accounting may add operational energy and allocated embodied energy only after checking that their scopes do not overlap. Let

$$E_{\text{lifecycle}} = E_{\text{operational}} + E_{\text{embodied,alloc}}.$$

The allocation may depend on lifetime, utilization, task count, or revenue. Those choices can materially change a per-task denominator. ISO lifecycle standards require goal and scope definition, inventory, impact assessment, and interpretation [16, 17]. A graph can store these bridges, but it does not choose them.

8 The informational layers: symbols, uncertainty, and computation

8.1 Shannon information

For a discrete random variable X with mass function $p(x)$, Shannon entropy is

$$H(X) = - \sum_x p(x) \log_2 p(x).$$

For jointly distributed X, Y , mutual information is

$$I(X; Y) = \sum_{x,y} p(x, y) \log_2 \frac{p(x, y)}{p(x)p(y)} = H(X) - H(X | Y).$$

These quantities describe uncertainty and statistical dependence under a probability model [29]. Their unit is the bit when the logarithm has base two.

A bits-per-joule ratio can be meaningful for a declared communication or storage task:

$$\eta_{\text{info}} = \frac{\text{reliably communicated information bits}}{E_{\text{system}}}.$$

The numerator needs a source distribution, channel model, code, error criterion, and time interval. It is not the number of characters in a message and not its economic significance.

Shannon explicitly bracketed semantics from the engineering theory. A rare symbol can carry high surprisal and still be false, useless, or harmful. A short command can have large consequences and low Shannon length. Meaning and value need later typed stages.

8.2 Landauer's principle

Landauer's result concerns logically irreversible operations implemented by a physical system in contact with a thermal environment [21]. In the idealized quasistatic limit, erasing one unknown bit requires at least

$$Q_{\min} = k_B T \ln 2$$

of heat dissipation to the environment. At $T = 300$ K, the scale is approximately 2.9×10^{-21} J per erased bit.

Four qualifications matter.

1. The bound concerns logically irreversible information loss, not every logical operation or every stored bit.
2. It is a lower bound under a physical model, not a prediction of the energy used by a contemporary computer.
3. Real systems spend energy on switching, memory, communication, leakage, clocks, control, error correction, cooling, and idle capacity.
4. The bound contains no task correctness, semantic usefulness, intelligence, adoption, price, or welfare term.

Bennett’s reversible-computation work shows that a computation can, in principle, be arranged so that logical steps are reversible, while input preparation, output handling, noise, and practical speed introduce other costs [7, 8]. It is therefore wrong to multiply a token count by $k_B T \ln 2$ and call the result the minimum energy of intelligence. Tokens are not erased bits, and intelligence has not been physically defined by that multiplication.

8.3 Prediction, semantics, and agents

Predictive information can be tied to thermodynamic efficiency in a specified stochastic process [31]. Semantic-information proposals such as Kolchinsky and Wolpert [18] ask which correlations contribute to an agent’s ability to maintain itself. These are valuable bridges because they state the agent and viability function. They also illustrate the general rule: semantics enters through additional structure.

For machine intelligence, a practical chain might be

operations $\rightarrow$ model output $\rightarrow$ scored task result $\rightarrow$ reliable task service.

The scoring rule can use accuracy, calibration, robustness, latency, or human evaluation. None is supplied by Shannon entropy alone.

8.4 Information is not a conserved economic input

Information can be copied at low marginal energy, combined with other information, forgotten, monopolized, or made obsolete. Its economic value depends on rights, timing, complements, and actions. A production function can include an information input, but that modeling choice does not make information conserved like energy.

Limitation 8.1 (Landauer scope). Landauer’s bound cannot rank economic systems by intelligence per joule unless the systems are first reduced to the same specified irreversible logical operation. That reduction would omit most of the economic question.

9 From verified outcomes to economic value

9.1 Tasks and usefulness

An immediate output becomes a verified outcome only under a task definition. Let O_i be output i , $S_i(O_i)$ a scoring rule, and z_i a success threshold. A successful task-equivalent contribution can be written

$$U_i = w_i \mathbf{1}\{S_i(O_i) \geq z_i\},$$

where w_i is a published difficulty or importance weight. The aggregate $\sum_i U_i$ is distribution-specific. Changing the task mix or weights can change a ranking.

The type of U_i includes model version, task distribution, scorer, threshold, latency constraint, and reliability window. Calling it “intelligence” does not remove those dependencies. The responsible label is “verified task equivalents under protocol P .”

9.2 Adoption

A correct output has no realized effect if nobody uses it. Let D_i indicate adoption and let $A_i = D_i U_i$. Adoption can depend on confidence, workflow design, incentives, legal authority, and user skill. It can also be selective: people may adopt outputs when private information suggests high benefit. Consequently, the average outcome among adopted recommendations does not identify the effect of making recommendations available.

An institution enters here as a rule system that changes which actions are authorized, recorded, or enforced. The institution may use physical energy through buildings and computing, but institutional order itself is not a thermodynamic state variable. Its operation should be modeled through rules, rights, information, incentives, and observed actions.

9.3 Counterfactual value

Let $Y_i(1)$ be an outcome for unit i under the deployed system and $Y_i(0)$ the outcome under a named counterfactual. Let W map the outcome vector to a monetary, welfare, capability, or domain-specific value. The incremental value is

$$\Delta W_i = W(Y_i(1)) - W(Y_i(0)).$$

Only one potential outcome is normally observed. Identification therefore requires randomization, a credible quasi-experiment, structural assumptions, or an explicitly conditional scenario.

The energy numerator and value denominator must refer to compatible populations and intervals. If energy is measured during a one-hour benchmark but value is projected over ten years, the model needs a bridge from the benchmark to deployment volume, reliability, maintenance, and discounting.

9.4 Prices, revenue, surplus, and welfare

These objects should not be conflated.

- Revenue is price times sales and can include transfers.
- Producer surplus subtracts variable opportunity costs under a stated market model.
- Consumer surplus depends on demand and prices.
- Profit follows an accounting convention and time horizon.
- Social welfare may add externalities and distributional weights.
- Capability or mission outcomes may remain vector-valued.

A ratio of market capitalization to annual electricity use can be computed, but it is not a stage yield. The numerator is a forward-looking asset valuation and the denominator is a flow over a chosen year. The graph should mark the comparison as **OBSTRUCTED** if it is presented as causal production.

9.5 A causal graph behind the accounting graph

The transduction graph records measured stages. A causal graph records assumed data-generating relations. They need not coincide. For example,

$$E \leftarrow L \rightarrow U$$

may arise because harder workloads L consume more energy and have different success rates. Regressing U on E can then misread workload composition as an energy effect. The accounting identity $U/E = (C/E)(O/C)(U/O)$ still holds for matched totals, but it does not identify the causal return to added energy.

Remark 9.1 (Conditional economic stage). Every edge into ΔW is **CONDITIONAL** unless the counterfactual is identified under a stated design. A deterministic accounting product does not upgrade that status.

10 Uncertainty and statistical decomposition

10.1 Measurement error

Each interface amount may have calibration error, sampling error, model error, or allocation uncertainty. Let the reported amount be

$$\tilde{X}_k = X_k + \epsilon_k.$$

If adjacent stages use separate instruments, then $\tilde{X}_k^{\text{out}}$ and $\tilde{X}_k^{\text{in}}$ may differ even when they target the same underlying quantity. Reconciliation should report the residual and its uncertainty.

Ratios can be unstable when the denominator is small. A confidence interval formed by treating numerator and denominator as independent may be wrong if they share calibration or load variation. Bootstrap or Bayesian models should resample at the experimental unit that preserves dependence.

10.2 Logarithmic decomposition

For positive yields,

$$\log Y_{\text{total}} = \sum_{k=1}^n \log Y_k.$$

This form is useful for multiplicative uncertainty. If log yields have covariance matrix Σ , then

$$\text{Var}(\log Y_{\text{total}}) = \mathbf{1}^\top \Sigma \mathbf{1}.$$

The off-diagonal terms matter. Assuming independence sets them to zero and should be labeled as an assumption.

10.3 Monte Carlo model

The reference implementation accepts a coefficient of variation for each stage and samples independent lognormal multiplicative factors with arithmetic mean one. For stage k , if the coefficient of variation is c_k ,

$$\sigma_k^2 = \log(1 + c_k^2), \quad \mu_k = -\frac{1}{2}\sigma_k^2.$$

The random factor $\exp(\mu_k + \sigma_k Z_k)$ then has mean one. This model keeps sampled yields non-negative and makes the assumption inspectable.

Independence is not a default truth. The implementation is a compact reference fixture. Empirical work should estimate dependence, include shared calibration variables, and run sensitivity cases.

10.4 Zero and negative outcomes

A stage can produce zero verified outcomes. Its path product is then zero. Negative economic value requires a signed terminal stage. A logarithmic decomposition does not apply to zero or negative values. One option is to maintain separate non-negative benefit and harm paths:

$$\Delta W = W^+ - W^-.$$

Another is to work with signed levels and avoid lognormal uncertainty. The choice affects summaries and must appear in the manifest.

10.5 Selection and missingness

Failed runs, abandoned tasks, and unadopted recommendations belong in the denominator and outcome distribution when they fall inside the boundary. Dropping them conditions on success and raises the apparent yield. Missing outcomes require an explicit missingness model or bounds. A graph with a “missing” branch makes this visible.

Limitation 10.1 (Uncertainty does not repair a bad boundary). A narrow confidence interval around a mismatched or causally uninterpretable ratio does not make that ratio meaningful. Boundary and identification checks come before precision.

11 Bounded case study I: useful hydraulic work

This section uses a constructed engineering example. It is not an empirical claim about a particular facility.

11.1 Boundary and stages

Consider a motor, pump, and pipe segment operating during one metered run. The input is 10.0 MJ of wall electricity. The motor delivers shaft work at an efficiency of 0.92. The pump converts shaft work to hydraulic work at 0.76. The pipe segment delivers 0.85 of that hydraulic work at the stated outlet head.

The typed chain is

$$E_{\text{wall}} \xrightarrow{0.92} W_{\text{shaft}} \xrightarrow{0.76} W_{\text{hydraulic}} \xrightarrow{0.85} W_{\text{outlet}}.$$

All four nodes use joules, but their semantic types identify different locations and functions.

Interface	Amount	Type	Boundary
Wall electricity	10.000 MJ	electrical energy	meter terminals
Shaft work	9.200 MJ	mechanical work	motor shaft
Hydraulic work	6.992 MJ	fluid work	pump outlet
Delivered work	5.9432 MJ	hydraulic work	pipe endpoint

Table 1: Constructed amounts for the hydraulic example.

The product is

$$0.92 \times 0.76 \times 0.85 = 0.59432.$$

The direct ratio is

$$\frac{5.9432}{10.000} = 0.59432.$$

This is an exact instance of Theorem 5.1.

11.2 Exergy interpretation

Under common engineering conditions, wall electricity is treated as pure exergy. The shaft and hydraulic work are also work forms. Exergy destruction and losses can be assigned to the motor, pump, and pipe model, subject to the reference environment and instrumentation. The calculation remains physical through the delivered-work node.

Suppose the service is “one cubic meter of potable water delivered at a specified pressure.” Converting delivered hydraulic work to that service requires flow, head, leakage, and water-quality measurements. Converting the service to economic value requires still more: scarcity, alternative supply, health outcomes, tariffs, distribution, and a counterfactual. The 0.59432 physical yield says none of those things.

11.3 What is proved and what is not

- The arithmetic composition is PROVED for the stated numbers.
- Reproduction by code is COMPUTATIONAL.
- The component efficiencies are assumed fixture inputs, not measured evidence.
- Any monetary value of delivered water is OPEN in this example.

The case shows why the framework does not need to deny physical efficiency. It preserves a clean physical chain and stops before an unsupported economic edge.

12 Bounded case study II: a task service

This case mirrors the executable fixture in `src/joule_standard/transduction.py`. The numbers are deliberately small and constructed so that every interface can be inspected. They do not describe a deployed model.

12.1 Declared chain

During one metered interval:

$$100 \text{ J} \rightarrow 1000 \text{ named operations} \rightarrow 25 \text{ verified task equivalents} \rightarrow 50 \text{ USD}_{2026}$$

under a hypothetical randomized ninety-day value study.

The stage yields are

$$y_1 = 10 \frac{\text{operation}}{\text{J}}, \quad y_2 = 0.025 \frac{\text{task}}{\text{operation}}, \quad y_3 = 2 \frac{\text{USD}_{2026}}{\text{task}}.$$

Their product is

$$y_1 y_2 y_3 = 0.5 \frac{\text{USD}_{2026}}{\text{J}},$$

which matches 50/100.

12.2 Why each adjective matters

“Named operations” means a stable hardware and software counting rule. Floating-point operations, instructions, and accelerator events cannot be mixed. “Verified task equivalents” means a fixed task distribution, scoring rule, threshold, and weighting scheme. “Incremental value” means the difference relative to a control condition under a declared study. Currency year and discounting must be stated.

If the final 50 dollars were merely revenue observed after deployment, the last edge would not have the claimed type. The product could still be computed as revenue per joule, but it would not be incremental value per joule.

12.3 Monte Carlo illustration

Assign coefficients of variation 0.05, 0.10, and 0.20 to the three stage yields. Under independent mean-one lognormal factors, the total coefficient of variation is

$$[(1 + 0.05^2)(1 + 0.10^2)(1 + 0.20^2) - 1]^{1/2} \approx 0.230.$$

The total mean remains 0.5 USD/J under the model, while the median is lower because the distribution is right-skewed. A Monte Carlo run with a fixed seed provides reproducible sample quantiles.

This result is COMPUTATIONAL. The lognormal form and independence are assumptions. A real experiment would estimate covariance from repeated runs and carry uncertainty in the counterfactual effect.

12.4 A rank reversal caused by the task mix

Suppose system A solves easy tasks at 0.9 success and hard tasks at 0.4, while system B obtains 0.7 and 0.6. If both use the same energy per attempt, A ranks higher when easy tasks dominate sufficiently, and B ranks higher when hard tasks receive sufficient weight. No hardware change is needed. The verified-outcome node changed because its task distribution changed.

This is not a defect in task evaluation. It is evidence that “intelligence” must remain distribution-specific. A universal intelligence-per-joule scalar is OBSTRUCTED without a justified task and weighting distribution.

13 Bounded case study III: a communication channel

Consider a binary symmetric channel with crossover probability $p = 0.01$ and uniform independent inputs. The capacity per channel use is

$$C = 1 - H_2(p),$$

where

$$H_2(p) = -p \log_2 p - (1 - p) \log_2 (1 - p).$$

Numerically, $C \approx 0.9192$ bits per use.

Suppose a measured system uses 1 mJ for one million channel uses, including transmitter, receiver, and coding within the declared boundary. At capacity in the asymptotic coding model, the ratio is approximately

$$\frac{0.9192 \times 10^6 \text{ bit}}{10^{-3} \text{ J}} = 9.192 \times 10^8 \frac{\text{bit}}{\text{J}}.$$

13.1 Scope of the result

The number describes reliable communication under a source and channel model. It does not depend on whether the bits encode a weather warning, random noise, or a fraudulent instruction. The semantic and economic consequences can differ while the communication ratio stays fixed.

To extend the chain, one might add

decoded bits $\rightarrow$ correctly interpreted warning $\rightarrow$ protective action $\rightarrow$ loss avoided.

Each arrow needs evidence. Language, user attention, institutional authority, and local hazards enter after channel decoding. Treating the original bits-per-joule figure as “meaning per joule” would skip those stages.

13.2 Relation to Landauer

The example’s 10^{-9} J per channel use is many orders of magnitude above $k_B T \ln 2$ at room temperature. That comparison does not identify avoidable inefficiency because the measured system performs communication, coding, control, and physical signaling rather than one ideal bit erasure. A useful engineering study would decompose the hardware energy and identify actual irreversible operations.

13.3 An adversarial message

Now hold the bit distribution, channel, code, and energy fixed while changing the decoded content from an accurate evacuation route to a false one. Shannon entropy and bits per joule are unchanged. Expected welfare can change sign. This construction is enough to refute a universal conversion from Shannon information to economic value.

Proposition 13.1 (Semantic nonidentification). *Channel statistics and energy alone do not identify task usefulness or economic value.*

Proof. Construct two messages with the same length and symbol probabilities, sent through the same channel with the same energy. Let one induce a beneficial action and the other an equally credible harmful action. The physical and Shannon quantities are identical while the task and value outcomes differ. $\square$

The proposition is PROVED as a nonidentification counterexample. The magnitude of any real effect remains empirical.

14 Institutions, coordination, and order

The motivating program includes coordination and institutional order among the outputs of energy-using systems. These terms need special care because metaphor can outrun measurement.

14.1 An institution is a rule system

Following North [24], institutions can be understood as formal and informal constraints that structure interaction. A registry, exchange, court, or standards body uses people, buildings, communication, and computing. Those physical operations have energy accounts. The institution also assigns rights, recognizes records, and enforces procedures. Those legal and social facts are not reducible to the heat dissipated by the hardware.

A typed path for a registry might be

$$E_{\text{facility}} \rightarrow \text{processed filings} \rightarrow \text{validated records} \rightarrow \text{recognized rights} \rightarrow \text{disputes avoided}.$$

Processed filings are counted events. Validation depends on rules and error audits. Recognized rights depend on legal authority. Disputes avoided are counterfactual. Only the first stage is a direct energy conversion.

14.2 Coordination measures

Coordination can be measured through domain outcomes such as wait time, conflict rate, matching quality, schedule adherence, or losses avoided. No one measure exhausts institutional order. A vector is usually more honest:

$$\mathbf{U}_{\text{coord}} = (\text{latency}, \text{error}, \text{access}, \text{appeal}, \text{distributional incidence}).$$

Aggregation requires weights or a constrained optimization problem.

14.3 Human labor and energy

Human metabolism uses energy, but converting labor compensation or skill into food calories is not a useful economic production identity. Labor services depend on knowledge, law, health, preferences, and institutions. A lifecycle study may count the energy of commuting, buildings, or equipment under a declared scope. It should not claim that human contribution is explained by metabolic joules.

The same warning applies to phrases such as “institutional energy.” Unless the variable denotes a physical energy flow, the phrase should be avoided. The graph uses institutional nodes with their own units and types.

14.4 Trust and assurance

Trust is a psychological or relational state. Assurance can sometimes be operationalized through audit coverage, failure probability, attack cost, settlement delay, or expected loss.

A later paper in this program treats proof-of-work settlement under a threat model. The present paper supplies only the type discipline: energy, computation, protocol events, scenario-conditioned assurance, and expected loss reduction are separate nodes.

Limitation 14.1 (No institutional state equation). There is no equation of state that maps energy and entropy to legitimacy, authority, or trust. Any such mapping requires behavioral and institutional assumptions outside thermodynamics.

15 Measurement protocol and data requirements

15.1 Pre-registration of the graph

Before collecting data, an empirical study should publish:

1. the graph nodes and their semantic types;
2. the edge definitions and observation procedures;
3. the physical and organizational boundary;
4. the functional unit and task distribution;
5. the counterfactual and estimand;
6. shared-infrastructure allocations;
7. the uncertainty and missingness model;
8. stopping rules and exclusions.

Pre-registration prevents the graph from being redrawn after outcomes are known. It also makes clear which stages are observed, modeled, or assumed.

15.2 Energy instrumentation

For computing systems, wall measurement is often the most defensible operational denominator because it can include host processors, memory, networking, storage, and power conversion. Facility measurement adds cooling and common infrastructure. Device telemetry can help diagnose components but should not be relabeled as wall energy. MLCommons power methodology offers a useful benchmark-specific precedent [23].

Sampling rate must capture workload transients. The study should separate warm-up, idle, failed requests, retries, and cooldown according to its deployment boundary. Power integration should report calibration and clock alignment.

15.3 Outcome instrumentation

Immediate outputs need stable counting rules. Verified outcomes need blinded or automated scoring, inter-rater reliability when humans score, and a policy for partial credit. Adoption needs logs or surveys tied to the same units. Economic outcomes need administrative or experimental data, with privacy and selection risks addressed.

15.4 Versioning and drift

Hardware firmware, compilers, model weights, prompts, task distributions, and institutional rules change. Each evidence record should include versions and dates. A yield estimated under one stack should not be transferred to another without validation.

15.5 Comparability gate

Two terminal-to-energy results are directly comparable only if:

- their functional units and terminal types match;
- their energy stages match or have a validated bridge;

- their time horizons and geographic contexts match;
- their counterfactuals and value functionals match;
- their allocations of shared infrastructure are compatible;
- uncertainty intervals and exclusions are available.

If one condition fails, the comparison is **OBSTRUCTED** or reported as a sensitivity range. A point ranking should not be forced.

16 Reference implementation

The companion module `src/joule_standard/transduction.py` implements the paper’s accounting core using the Python standard library.

16.1 Data model

- `Unit` stores symbolic unit exponents and supports multiplication, division, powers, and cancellation.
- `QuantityType` stores a semantic name, unit, and boundary.
- `Node` assigns a unique interface name to a quantity type.
- `Stage` records source, target, input amount, and output amount.
- `TransductionGraph` rejects missing endpoints and cycles.
- `deterministic_decomposition` verifies path and interface equality before multiplying yields.
- `monte_carlo_decomposition` propagates declared independent lognormal yield uncertainty with a fixed seed.

The implementation keeps exergy and energy distinct even when both use the unit J. It keeps verified tasks and operations distinct even when both are represented as counts. A path can compose only through the same named node.

16.2 Verification contract

The tests check:

1. dimensional cancellation;
2. semantic distinction under a common physical unit;
3. exact direct-versus-product equality;
4. rejection of unmatched nodes and interface amounts;
5. cycle rejection without partial mutation;
6. deterministic seeded Monte Carlo output;
7. convergence of the sample mean to the declared model mean;
8. rejection of invalid uncertainty inputs;
9. read-only graph views.

Passing tests support the COMPUTATIONAL claim that the reference implementation enforces these fixtures. They do not prove that external data have correct boundaries or that a causal estimand is identified.

16.3 Why the library is small

The library does not include a unit conversion database, probability programming system, lifecycle inventory, or causal estimator. Those belong to specialized tools. The reference code demonstrates the minimum contract that future adapters must preserve.

16.4 Determinism and reproducibility

Monte Carlo sampling uses a local pseudorandom number generator and explicit seed. The result object includes all draws, summary statistics, path, unit, sample count, and seed. A production implementation may avoid retaining every draw, but it should retain enough provenance to reproduce quantiles.

17 Falsifiable claims and evidence status

Table 2 separates mathematical results, computational evidence, conditional empirical claims, and open questions.

Status	Claim	Falsifier or boundary
PROVED	Matched stage yields telescope to the direct terminal-to-initial ratio.	An interface differs in type or amount.
PROVED	Unit rescaling changes local numerical yields covariantly and preserves the represented path ratio.	A conversion changes the underlying quantity, not only its unit.
PROVED	Energy and channel statistics alone do not identify semantic or economic value.	A task or value functional is added, in which case the claim no longer says “alone.”
COMPUTATIONAL	The Python fixtures enforce unit and interface checks and reproduce deterministic and seeded Monte Carlo decompositions.	A test fails or an accepted invalid path is constructed.
CONDITIONAL	A measured task service has incremental value per joule.	The counterfactual, value functional, or causal identification fails.
OBSTRUCTED	A universal scalar converts joules to intelligence, institutional order, or value without task and normative structure.	A domain-independent conversion is derived without added structure.
OPEN	Typed reporting improves empirical decisions, reproducibility, or policy.	Requires comparative use in real studies.

Table 2: Claim-status matrix.

17.1 Direct empirical tests

The framework can be tested without accepting its broader research program. Researchers can ask:

1. Do independent analysts reproduce a reported path from its manifest?
2. How often do published comparisons fail the comparability gate?
3. Does stage decomposition predict where interventions improve terminal outcomes?
4. Do covariance-aware uncertainty models materially change rankings?
5. Do task and value distributions produce anticipated rank reversals?
6. Does explicit attribution reduce double counting in joint production?

Negative answers would narrow the framework’s practical value even though the algebraic theorem remains true.

17.2 A proposed adversarial audit

An audit should attempt five attacks:

1. substitute device energy for facility energy;
2. substitute output count for verified outcome count;
3. drop failed or unadopted cases;
4. count one shared input in full on two branches;
5. replace a counterfactual effect with observed revenue.

If the manifest or software accepts these substitutions without warning, the implementation has failed its intended purpose.

18 Limitations

18.1 The main theorem is an identity

The stage-yield theorem does not discover a behavioral law. Its scientific content lies in the preconditions and evidence records. A critic may reasonably find the identity too elementary. The empirical response must be to show that enforcing its interfaces prevents errors or improves decisions. That evidence does not yet exist.

18.2 Types can become bureaucratic

A type system can produce elaborate metadata without better measurement. Boundaries may be vague, copied, or strategically chosen. The framework needs independent audits and concise manifests, not merely more fields.

18.3 No universal ontology

Different fields use different concepts of useful work, task success, assurance, and welfare. The graph does not impose a complete ontology. This flexibility supports application but can weaken cross-study comparability. Registries and domain-specific profiles will be needed.

18.4 No equilibrium or behavioral theory

The graph records flows and conditional relations. It does not determine prices, demand, strategic behavior, market power, rebound, or general equilibrium. Those require separate models. Later papers in this program study marginal allocation and rebound under explicit assumptions.

18.5 No causal identification from composition

Exact multiplication of observed yields does not identify the effect of increasing energy or improving one stage. Confounding, selection, spillovers, and equilibrium responses remain. The causal graph and research design must be evaluated separately.

18.6 Uncertainty model is illustrative

Independent lognormal stage factors are convenient for positive quantities. They can be wrong for bounded success rates, zero-inflated outcomes, heavy tails, or correlated stages. The code is not a general uncertainty engine.

18.7 Economic value remains contested

Prices, profits, welfare, capabilities, and mission outcomes answer different questions. Distributional weights and rights cannot be derived from joules. Publishing a vector does not eliminate normative choice; it postpones aggregation and makes the choice visible.

18.8 Lifecycle data are difficult

Embodied energy, equipment lifetime, utilization, supply-chain location, and marginal grid effects are uncertain. A lifecycle denominator can appear more comprehensive while importing large allocation uncertainty. Operational and lifecycle results should often be reported side by side.

18.9 Metaphor remains a risk

The word “transduction” comes from physical and engineering contexts. Here it means a typed sequence of measured or modeled transformations. It does not claim that institutions, intelligence, or value are physical energy forms. Information, intelligence, institutions, and value are not thermodynamic state variables.

19 Discussion

The framework changes the usual question. Instead of asking how many dollars, tokens, or units of trust a joule “becomes,” it asks which typed interfaces connect an energy measurement to a later outcome. That shift makes it possible to compare the structure of systems without declaring their outputs identical.

The physical part of a chain can use conservation and exergy balances. The information part can use channel models and error criteria. The task part can use benchmark or field evaluation. The institutional part can use rules, records, and behavioral data. The economic part can use causal and welfare analysis. The graph is common, but the evidence at each edge remains discipline-specific.

This separation also clarifies intervention. Improving motor efficiency addresses a physical stage. Improving a model’s scoring reliability addresses an outcome stage. Changing workflow incentives addresses adoption. Changing a market or legal rule addresses realization. A high upstream yield does not guarantee a high terminal yield.

The same logic cautions against thermodynamic rhetoric in economics. Energy is indispensable to material and digital production, but scarcity, preference, rights, and coordination shape value. The entropy law constrains physical processes. It does not by itself choose a social welfare function. Conversely, economic prices do not repeal physical constraints. A typed graph allows both statements to remain true.

19.1 Implications for AI measurement

Tokens per joule can be useful for capacity planning under a fixed tokenizer, model, output policy, and quality target. It is not a measure of intelligence or economic value. A task service should report at least wall energy, task distribution, correctness, reliability, latency, failure rate, and uncertainty. Economic claims need adoption and counterfactual outcomes.

19.2 Implications for national accounting

At national scale, primary energy, final energy, useful work, digital services, and value added should appear in linked but nonduplicative accounts. Adding intermediate and final flows creates double counting. A satellite account can publish vectors and bridge tables without replacing gross domestic product.

19.3 Implications for market allocation

Average value per joule does not answer where the next joule should go. Marginal allocation depends on time, location, congestion, reliability, externalities, and workload constraints. The graph can supply local response functions, but optimization and equilibrium require later analysis.

20 Conclusion

Energy-using systems can be studied within one accounting language without pretending that all outputs share a physical essence. The language is a typed transduction graph. Nodes identify quantities, units, boundaries, intervals, counterfactuals, and attribution rules. Stages record measured or modeled relations. A path composes only when its interfaces match.

Under those conditions, stage yields telescope. The theorem is exact and limited. It localizes shortfalls and makes missing evidence visible. It does not turn energy into exergy, exergy into information, information into intelligence, or intelligence into value. Those transitions require physical, statistical, task, institutional, causal, and normative definitions.

Exergy remains maximum useful work relative to an environment. Useful work remains a physical or service quantity under a conversion convention. Shannon information remains uncertainty reduction under a probability model. Landauer's bound remains a constraint on logically irreversible operations under a thermal model. Economic value remains relative to a counterfactual and value functional.

The practical test now lies ahead. Researchers should use the graph on real systems, publish its manifests, invite adversarial boundary checks, and see whether the decomposition improves intervention and comparison. Until then, the mathematical claims are **PROVED**, the software evidence is **COMPUTATIONAL**, empirical value claims are **CONDITIONAL**, and a universal joule-to-intelligence conversion is **OBSTRUCTED**.

A Formal type algebra

Let G be the free abelian group over a set of unit symbols. A unit is a finite map from symbols to integer exponents. Multiplication adds exponent vectors; division subtracts them. The identity has every exponent zero.

A semantic quantity type is not an element of G . It is a record

$$\theta = (q, d, m),$$

where q is a semantic identifier, $d \in G$, and m is boundary and measurement metadata. The projection $\dim(\theta) = d$ forgets semantics. Therefore

$$\theta_1 = \theta_2 \implies \dim(\theta_1) = \dim(\theta_2),$$

but the converse does not hold.

For a stage $\ell : \theta_i \rightarrow \theta_j$, its yield has formal unit

$$\dim(y_\ell) = \dim(\theta_j) \dim(\theta_i)^{-1}.$$

For a composable path, multiplication gives

$$\prod_{k=1}^n \dim(\theta_k) \dim(\theta_{k-1})^{-1} = \dim(\theta_n) \dim(\theta_0)^{-1}.$$

This unit cancellation mirrors the numerical proof.

A.1 Typed versus untyped equality

Suppose $E = 100 \text{ J}$ is electrical energy and $B = 100 \text{ J}$ is exergy under a reference environment. Their numerical magnitudes and units match, but their semantic types differ. The statement $E = B$ needs a physical argument about the carrier and reference environment. The type system does not infer it from formatting.

Likewise, two counts of 25 may be operations and verified tasks. An evaluation stage, not arithmetic equality, connects them.

B Additional proofs and counterexamples

B.1 Boundary reversal

Let systems A and B each deliver 100 task equivalents. Their device energies are 50 and 60 joules, so A ranks above B under the device boundary. Suppose A requires 40 joules of host and facility overhead while B requires 10. Facility totals are 90 and 70, so B ranks above A. The ranking changed because the boundary changed.

This counterexample is PROVED and does not show which boundary is correct. The application chooses the boundary.

B.2 Allocation reversal

Suppose two services share 100 joules of idle infrastructure. Under equal allocation, each receives 50 joules. Under allocation by peak capacity, one receives 80 and the other 20. If their direct energies and outputs are close, the efficiency ranking can reverse. No arithmetic error is needed. The allocation rule belongs in the type.

B.3 Expectation counterexample

Let Y_1 be equally likely to be 1 or 2, and set $Y_2 = 1/Y_1$. Then $Y_1 Y_2 = 1$ surely, so

$$\mathbb{E}[Y_1 Y_2] = 1.$$

But

$$\mathbb{E}[Y_1] \mathbb{E}[Y_2] = \frac{3}{2} \frac{3}{4} = \frac{9}{8}.$$

Multiplying mean stage yields overstates the mean total yield. Dependence matters.

B.4 Duplicate-branch counterexample

One inference run uses 100 joules and produces a recommendation copied to two departments. Each department records the full 100 joules and a 10-dollar benefit. If the branches are summed, the study reports $20/200 = 0.1$ dollars per joule, which happens to equal the correctly combined $20/100 = 0.2$ only if the denominator is handled differently. Charging full energy to each branch and then interpreting branch values as additive gives an incoherent total. An attribution-complete graph records one shared input and two outcome branches.

C Reporting manifest

A machine-readable or tabular manifest should include the following fields.

1. Study identifier and version.
2. Node identifier and semantic quantity name.
3. Unit and currency base year when relevant.
4. Physical and organizational boundary.
5. Time interval and geographic location.
6. Counterfactual.
7. Value or outcome functional.
8. Energy convention.
9. Attribution and shared-infrastructure rule.
10. Observation or model source.
11. Calibration and validation evidence.
12. Uncertainty distribution and dependence assumptions.
13. Missingness and failure treatment.
14. Stage source and target.
15. Reconciliation residual.

16. Evidence status.
17. Exclusions and known limitations.

D Glossary

Energy A physical state-function quantity governed by the first law, reported under a specified system boundary.

Exergy Maximum useful work relative to a specified environment as a system approaches equilibrium with that environment.

Useful work Delivered physical work or an explicitly constructed end-use work aggregate.

Shannon information A measure of uncertainty or statistical dependence under a probability model.

Landauer bound A thermal lower bound associated with logically irreversible information erasure under stated conditions.

Verified task equivalent An outcome count under a fixed task distribution, scorer, threshold, and weighting rule.

Adoption Observed use of an output in a specified workflow.

Incremental value Difference in a declared value functional between treatment and counterfactual outcomes.

Institution A rule system that structures interaction, rights, and enforcement. It is not a thermodynamic state.

Stage yield Output amount divided by input amount for one typed edge.

Composable path A sequence whose adjacent type, boundary, and amount interfaces match.

References

- [1] D. Acemoglu, V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi, “The network origins of aggregate fluctuations,” *Econometrica*, vol. 80, no. 5, pp. 1977 to 2016, 2012. <https://doi.org/10.3982/ECTA9623>
- [2] K. J. Arrow, *Social Choice and Individual Values*. New York: Wiley, 1951.
- [3] R. U. Ayres and B. Warr, “Accounting for growth: the role of physical work,” *Structural Change and Economic Dynamics*, vol. 16, no. 2, pp. 181 to 209, 2005. <https://doi.org/10.1016/j.strueco.2003.10.003>
- [4] J. C. Baez and J. Master, “Open Petri nets,” *Mathematical Structures in Computer Science*, vol. 30, no. 3, pp. 314 to 341, 2020. <https://doi.org/10.1017/S0960129520000043>
- [5] J. C. Baez and B. S. Pollard, “A compositional framework for reaction networks,” *Reviews in Mathematical Physics*, vol. 29, no. 9, 1750028, 2017. <https://doi.org/10.1142/S0129055X17500283>
- [6] A. Bejan, G. Tsatsaronis, and M. Moran, *Thermal Design and Optimization*. New York: Wiley, 1996.
- [7] C. H. Bennett, “Logical reversibility of computation,” *IBM Journal of Research and Development*, vol. 17, no. 6, pp. 525 to 532, 1973. <https://doi.org/10.1147/rd.176.0525>
- [8] C. H. Bennett, “The thermodynamics of computation, a review,” *International Journal of Theoretical Physics*, vol. 21, pp. 905 to 940, 1982. <https://doi.org/10.1007/BF02084158>
- [9] V. M. Carvalho, “From micro to macro via production networks,” *Journal of Economic Perspectives*, vol. 28, no. 4, pp. 23 to 48, 2014. <https://doi.org/10.1257/jep.28.4.23>
- [10] C. J. Cleveland, R. Costanza, C. A. S. Hall, and R. Kaufmann, “Energy and the U.S. economy: a biophysical perspective,” *Science*, vol. 225, no. 4665, pp. 890 to 897, 1984. <https://doi.org/10.1126/science.225.4665.890>
- [11] B. Coecke, T. Fritz, and R. W. Spekkens, “A mathematical theory of resources,” *Information and Computation*, vol. 250, pp. 59 to 86, 2016. <https://doi.org/10.1016/j.ic.2016.02.008>
- [12] I. Dincer and M. A. Rosen, *Exergy: Energy, Environment and Sustainable Development*, 2nd ed. Oxford: Elsevier, 2013.
- [13] B. Fong and D. I. Spivak, *An Invitation to Applied Category Theory: Seven Sketches in Compositionality*. Cambridge: Cambridge University Press, 2019. <https://doi.org/10.1017/9781108668804>

- [14] C. R. Hulten, “Growth accounting with intermediate inputs,” *Review of Economic Studies*, vol. 45, no. 3, pp. 511 to 518, 1978. <https://doi.org/10.2307/2297252>
- [15] G. W. Imbens and D. B. Rubin, *Causal Inference for Statistics, Social, and Biomedical Sciences*. Cambridge: Cambridge University Press, 2015.
- [16] International Organization for Standardization, *ISO 14040:2006 Environmental Management, Life Cycle Assessment, Principles and Framework*. Geneva: ISO, 2006. <https://www.iso.org/standard/37456.html>
- [17] International Organization for Standardization, *ISO 14044:2006 Environmental Management, Life Cycle Assessment, Requirements and Guidelines*. Geneva: ISO, 2006. <https://www.iso.org/standard/38498.html>
- [18] A. Kolchinsky and D. H. Wolpert, “Semantic information, autonomous agency and non-equilibrium statistical physics,” *Interface Focus*, vol. 8, no. 6, 20180041, 2018. <https://doi.org/10.1098/rsfs.2018.0041>
- [19] T. J. Kotas, *The Exergy Method of Thermal Plant Analysis*. London: Butterworths, 1985.
- [20] R. Kümmel, *The Second Law of Economics: Energy, Entropy, and the Origins of Wealth*. New York: Springer, 2011. <https://doi.org/10.1007/978-1-4419-9365-6>
- [21] R. Landauer, “Irreversibility and heat generation in the computing process,” *IBM Journal of Research and Development*, vol. 5, no. 3, pp. 183 to 191, 1961. <https://doi.org/10.1147/rd.53.0183>
- [22] W. Leontief, *Input-Output Economics*, 2nd ed. New York: Oxford University Press, 1986.
- [23] MLCommons, “MLPerf inference power measurement methodology,” 2025. <https://docs.mlcommons.org/inference/power/>
- [24] D. C. North, *Institutions, Institutional Change and Economic Performance*. Cambridge: Cambridge University Press, 1990.
- [25] J. M. R. Parrondo, J. M. Horowitz, and T. Sagawa, “Thermodynamics of information,” *Nature Physics*, vol. 11, pp. 131 to 139, 2015. <https://doi.org/10.1038/nphys3230>
- [26] J. Pearl, *Causality: Models, Reasoning, and Inference*, 2nd ed. Cambridge: Cambridge University Press, 2009.
- [27] D. B. Rubin, “Estimating causal effects of treatments in randomized and nonrandomized studies,” *Journal of Educational Psychology*, vol. 66, no. 5, pp. 688 to 701, 1974. <https://doi.org/10.1037/h0037350>
- [28] A. Sen, “Personal utilities and public judgements: or what’s wrong with welfare economics?” *Economic Journal*, vol. 89, no. 355, pp. 537 to 558, 1979. <https://doi.org/10.2307/2231867>

- [29] C. E. Shannon, “A mathematical theory of communication,” *Bell System Technical Journal*, vol. 27, pp. 379 to 423 and 623 to 656, 1948. <https://doi.org/10.1002/j.1538-7305.1948.tb01338.x>
- [30] D. I. Stern, “The role of energy in economic growth,” *Annals of the New York Academy of Sciences*, vol. 1219, pp. 26 to 51, 2011. <https://doi.org/10.1111/j.1749-6632.2010.05921.x>
- [31] S. Still, D. A. Sivak, A. J. Bell, and G. E. Crooks, “Thermodynamics of prediction,” *Physical Review Letters*, vol. 109, 120604, 2012. <https://doi.org/10.1103/PhysRevLett.109.120604>
- [32] J. Szargut, D. R. Morris, and F. R. Steward, *Exergy Analysis of Thermal, Chemical, and Metallurgical Processes*. New York: Hemisphere, 1988.
- [33] G. Tsatsaronis, “Thermoeconomic analysis and optimization of energy systems,” *Progress in Energy and Combustion Science*, vol. 19, no. 3, pp. 227 to 257, 1993. [https://doi.org/10.1016/0360-1285\(93\)90016-8](https://doi.org/10.1016/0360-1285(93)90016-8)
- [34] A. Valero, M. A. Lozano, L. Serra, G. Tsatsaronis, J. Pisa, C. Frangopoulos, and M. R. von Spakovsky, “CGAM problem: definition and conventional solution,” *Energy*, vol. 19, no. 3, pp. 279 to 286, 1994. [https://doi.org/10.1016/0360-5442\(94\)90112-0](https://doi.org/10.1016/0360-5442(94)90112-0)
- [35] B. Warr, R. U. Ayres, N. Eisenmenger, F. Krausmann, and H. Schandl, “Energy use and economic development: a comparative analysis of useful work supply in Austria, Japan, the United Kingdom and the US during 100 years of economic growth,” *Ecological Economics*, vol. 69, no. 10, pp. 1904 to 1917, 2010. <https://doi.org/10.1016/j.ecolecon.2010.03.021>
- [36] B. Warr and R. U. Ayres, “Useful work and information as drivers of economic growth,” *Ecological Economics*, vol. 73, pp. 93 to 102, 2012. <https://doi.org/10.1016/j.ecolecon.2011.09.006>